Towards linking correctness conditions for concurrent objects and contextual trace refinement

Brijesh Dongol
Department of Computer Science, Brunel University London
Brijesh.Dongol@brunel.ac.uk

Lindsay Groves
School of Engineering and Computer Science, Victoria University of Wellington
lindsay@ecs.vuw.ac.nz

Correctness conditions for concurrent objects describe how atomicity of an abstract sequential object may be decomposed. Many different concurrent objects and proof methods for them have been developed. However, arguments about correctness are conducted with respect to an object in isolation. This is in contrast to real-world practice, where concurrent objects are often implemented as part of a programming language library (e.g., `java.util.concurrent`) and are instantiated within a client program. A natural question to ask, then is: How does a correctness condition for a concurrent object ensure correctness of a client program that uses the concurrent object? This paper presents the main issues that surround this question and provides some answers by linking different correctness conditions with a form of trace refinement.

1 Introduction

Concurrent objects provide operations that may be invoked by the parallel threads of a concurrent client, enabling more efficient computation in, for example, modern multi-core architectures. A concurrent object could be a commonly-used data structure such as a stack, queue, or set, or could implement a novel programming paradigm such as a software transactional memory object, which allows several reads and writes to be treated as a single atomic transaction [7].

Correctness conditions, such as linearizability [8], for concurrent objects cannot be defined in terms of pre/post conditions for their operations due to the possibility of interference while the operations are executing. Instead, they are defined in terms of a relation between (concurrent) histories of a concrete implementation and (sequential) histories of its abstract counterpart, which records occurrence of certain events at the concrete and abstract levels, such as invocations and responses of method calls [2].

Many papers have been devoted to verifying linearizability of concurrent object implementations [1]; however, these typically only consider behaviours of the concurrent object at hand in isolation — they do not provide any guarantees to the client programs that use concurrent objects. Therefore, letting $P[O]$ denote a client program P that uses object O, we consider the following question:

> *Provided concurrent object OC is correct with respect to sequential object OA, how are the behaviours of $P[OA]$ related to those of $P[OC]$?*

This question was been examined by Filipović et al. [3], who establish a link between two correctess conditions: *sequential consistency* and *linearizability* with a notion of refinement called *observational refinement*, which is essentially *data refinement* as defined by He et al. [6]. They show that sequential consistency and observational refinement coincide when threads are *data independent* (i.e., there is no

Eerke Boiten, John Derrick & Steve Reeves (Eds.):
Refinement Workshop 2015 (Refine '15)
EPTCS 209, 2016, pp. 107–111, doi:10.4204/EPTCS.209.8

```
Init:  Head = null

push(v) ==
H1: n := new(Node);
H2: n.val := v;
    repeat
H3:   ss := Head;
H4:   n.next := ss;
H5: until
      CAS(Head,ss,n)
H6: return
```

```
pop ==
    repeat
P1:   ss := Head;
P2:   if ss = null
P3:     then return empty
P4:   ssn := ss.next;
P5:   lv := ss.val
P6:  until
       CAS(Head,ss,ssn);
P7:  return lv
```

Figure 1: The Treiber stack

```
Init:  S = ⟨⟩

push(v) == atomic { S := ⟨v⟩⌢S }

pop ==
 atomic {
   if S = ⟨⟩ then
      return empty
   else lv := head(S) ;
      S := tail(S) ;
      return lv
 }
```

Figure 2: Abstract stack specification

data shared between client threads), while linearizability coincides with observational refinement when data dependence is allowed.

This paper provides a brief overview of our investigation, which examines this question in a more general setting than Filipović et al. A more detailed account of this work will be published elsewhere.

2 Concurrent objects and their clients

Figure 1 presents a simplified version of a non-blocking stack example due to Treiber [10], which has become a standard case study from the literature.[1] The implementation has fine-grained atomicity, and each line of the `push` and `pop` operations corresponds to a single atomic step. Synchronisation is achieved using an atomic compare-and-swap (`CAS`) instruction, which takes as input a *(shared) variable* `gv`, an *expected value* `lv` and a *new value* `nv`:

```
CAS(gv, lv, nv) ≙ atomic { if gv = lv then gv := nv ; return true
                           else return false }
```

The Treiber stack implements the abstract stack specification in Figure 2, where '⟨' and '⟩' delimit sequences, '⟨⟩' denotes the empty sequence, and '⌢' denotes sequence concatenation. The abstract stack consists of a sequence of elements S together with two operations *push* and *pop*. Note that when the stack is empty, *pop* returns a special value `empty` that cannot be pushed onto the stack.

A correctness condition is a relationship between the *histories* of the concrete and abstract systems. Each history records the interactions between a client and its objects. Typically, these are invocation and return events of operation calls, which form the object's external interface. Concurrent histories may consist of both overlapping and non-overlapping operation calls, inducing a partial order on events. Correctness conditions define how, if at all, this order is maintained in the corresponding abstract history. There are several well-known existing correctness conditions [7]. In this paper, we study two of these in detail: sequential consistency and linearizability.

- *Sequential consistency* [9] is a simple condition requiring the order of operation calls in a concrete history for a single process to be preserved. Operation calls performed by different processes may be reordered in the abstract history even if the operation calls do not overlap in the concrete history.

[1] We assume that garbage collection is used — this avoids the so-called ABA problem, where modifications to a shared pointer may go undetected when to the value changes from some value A to another value B then back to A.

- *Linearizability* [8] strengthens sequential consistency by requiring the order of non-overlapping operations to be preserved. Operation calls that overlap in the concrete history may be reordered when mapping to an abstract history.

Concurrent clients. Correctness conditions are usually defined in terms of a *most general client* which characterises the allowable behaviours of a concurrent object; however, they do not allow us to reason about specific clients that use these concurrent objects.

Example 1. The program below consists of threads 1 and 2, a shared stack s, and shared variables x, y and z. Thread 1 pushes 1 then 2 onto s, then pops the top element of s, and stores it in x. Concurrently, thread 2 pops the top element of s and stores it in y, then reads the value of x and stores it in z.

```
Init x, y, z = 0, 0, 0;
Thread 1:                                  Thread 2:
  T1:  s.push(1);                            U1:  y := s.pop();
  T2:  s.push(2);                            U2:  z := x;
  T3:  x := s.pop();
```

The program executes by *interleaving* the atomic statements of the two threads. In addition, depending on the implementation of s, we will get different behaviours of the client program because the effects of the concurrent operations on s may appear to occur in different orders. For example, s could be an instance of the Treiber Stack (Figure 1) (which is linearizable with respect to Figure 2), or some other stack that satisfies a different correctness condition (e.g., quiescent consistency [7]) with respect to the abstract stack in Figure 2.

With an example client program in place, we now return to the main question for this paper: How does one judge correctness of a system consisting of both a client and the objects it uses? More specifically, how does a correctness condition guaranteed by a concurrent object that a client uses affect the behaviour of the client itself? We will also consider what sorts of client behaviours should be examined. We address these issues as follows:

- First, we pin down the aspects of the system that are visible to an external observer. Following Filipović et al. [3], we take the observable state to be the state of the client variables, and the unobservable state to be the state of the objects they use. Therefore for the program in Figure 1, variables x, y and z are observable, but none of the variables of the stack implementation s are observable. This allows us to reason about a client with respect to different implementions of s.
- Second, we determine *when* a system may be observed. Unlike Filipović et al. [3], who only observe the state at the beginning and end of a client's execution, we take the states *throughout* a client's execution to be visible. This allows us to accommodate, for example, reactive or interactive clients, which may not terminate.

Taking both issues into account, our notion of correctness for the combined system will be a form of *contextual refinement*, which holds iff every (observable) trace of a client that uses a concurrent object is equivalent to some (observable) trace of the client using the abstract specification.

We say that *TS contextually trace refines AS with respect to the client program P* (denoted $AS \sqsubseteq_P TS$) iff every trace of $P[TS]$ is a possible trace of $P[AS]$. In this paper, we wish to know not only whether there is an abstract client trace equivalent to every concrete client trace, but also whether contextual refinement holds for every client program. To this end, we say *TS contextually trace refines AS* (denoted $AS \sqsubseteq TS$) iff *TS* contextually trace refines *AS* with respect to every client program *P*.

3 Linking correctness and contextual trace refinement

We now use the framework from the previous sections to explore the links between some well-known correctness conditions and contextual trace refinement.

Sequential consistency. Our main result for sequential consistency and contextual trace refinement is negative — sequential consistency does not guarantee contextual trace refinement of the underlying clients, regardless of whether the client program in question is data independent.

Lemma 2. *Suppose N is a client object and OA, OC are concurrent objects such that OC is sequentially consistent with respect to OA. Then it is not necessarily the case that $N[OA] \sqsubseteq N[OC]$ holds.*

Example 3. Consider the program below, where the client threads are data independent — `x` is local to thread `1`, while `y` and `z` are local to thread `2` — and s is assumed to be sequentially consistent.

```
Init x, y, z = 0;
Thread 1 ==                          Thread 2 ==
   T1: s.push(1);                       U1: z  := 1;
   T2: s.push(2);                       U2: out2  := s.pop();
   T3: out1  := s.pop();                U3: y  := out2;
   T4: x  := out1;
```

Suppose thread `1` is executed to completion, and then thread `2` is executed to completion. Because `s` is sequentially consistent, the first `pop` (at `T3`) may set $\mathtt{out}_1$ to `1`, the second (at `U2`) may set $\mathtt{out}_2$ to `2`. This gives the execution $\langle (x,y,z) \mapsto (0,0,0),\ \ (x,y,z) \mapsto (1,0,0),\ \ (x,y,z) \mapsto (1,0,1),\ \ (x,y,z) \mapsto (1,2,1)\rangle$, which cannot be generated when using the abstract stack *AS* from Figure 2 for `s`. □

Lemma 2 differs from the results of Filipović et al. [3], who show that for data independent clients, sequential consistency implies observational refinement. In essence, their result holds because observational refinement only considers the initial and final states of a client program — the intermediate states of a client's execution are ignored. Thus, internal reorderings due to sequentially consistent objects have no effect when only observing pre/post states. One can develop hiding conditions so that observational refinement is treated as a special case of contextual trace refinement, allowing one to obtain a positive result for sequential consistency equivalent to the result by Filipović et al. Full development of this theory is left for future work.

Linearizability. We now consider the link between linearizability and contextual trace refinement.

Example 4. Consider the program in Example 3, but now assume that the stack `s` is linearizable, e.g., is the Treiber stack. Reasoning that the traces generated by this program are valid (i.e., have a corresponding abstract trace) requires case analysis. In the final state we have either `x = 1` or `x = 2`. For traces that end with `x = 1`, `U2` must linearize before `T3`, but after `T2`. Therefore, `U1` also occurs before `T3`. For traces that end with `x = 1`, either `U2` linearizes before `T1` or `T3` linearizes before `U2`. The graph below depicts all possible traces of the client using the concrete implementation.

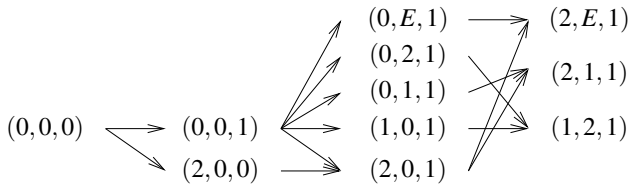

Each of these traces is also a possible trace of the client when it uses the abstract object. □

The next lemma states that when a concurrent object is linearizable with respect to an abstract object, then it also guarantees contextual trace refinement of clients that use it.

Lemma 5. *Suppose N is a client object, and OA and OC are concurrent objects such that OC is linearizable with respect to OA. Then $N[OA] \sqsubseteq N[OC]$ holds.*

4 Conclusions

In this paper, we have set up a framework for *studying the links between different correctness conditions for concurrent objects and contextual trace refinement*, which generalises the results of Filipović et al. [3]. We study sequential consistency and linearizability, and show that sequential consistency does not ensure contextual trace refinement. We have also shown that linearizability between an abstract specification and its linearizable implementation implies contextual trace refinement.

Gotsman and Yang [4] also extend Filipović et al.'s work, treating non-termination as abort, and considering both safety and progress properties (lock-freedom). Our trace-based framework can *distinguish* between non-terminating and aborting programs [5], and hence is more general than Gotsman and Yang [4]; though liveness properties are to be considered in future work.

References

[1] B. Dongol & J. Derrick (2015): *Verifying Linearisability: A Comparative Survey*. *ACM Comput. Surv.* 48(2), pp. 19:1–19:43, doi:10.1145/2796550.

[2] B. Dongol, J. Derrick, L. Groves & G. Smith (2015): *Defining Correctness Conditions for Concurrent Objects in Multicore Architectures*. In J. T. Boyland, editor: *ECOOP*, *LIPIcs* 37, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 470–494, doi:10.4230/LIPIcs.ECOOP.2015.470.

[3] I. Filipovic, P. W. O'Hearn, N. Rinetzky & H. Yang (2010): *Abstraction for concurrent objects*. *Theor. Comput. Sci.* 411(51-52), pp. 4379–4398, doi:10.1016/j.tcs.2010.09.021.

[4] A. Gotsman & H. Yang (2011): *Liveness-Preserving Atomicity Abstraction*. In L. Aceto, M. Henzinger & J. Sgall, editors: *ICALP(2)*, *LNCS* 6756, Springer, pp. 453–465, doi:10.1007/978-3-642-22012-8_36.

[5] I. J. Hayes, S. Dunne & L. Meinicke (2010): *Unifying Theories of Programming That Distinguish Nontermination and Abort*. In C. Bolduc, J. Desharnais & B. Ktari, editors: *MPC*, *LNCS* 6120, Springer, pp. 178–194, doi:10.1007/978-3-642-13321-3_12.

[6] J. He, C. A. R. Hoare & J. W. Sanders (1986): *Data Refinement Refined*. In B. Robinet & R. Wilhelm, editors: *ESOP*, *LNCS* 213, Springer, pp. 187–196, doi:10.1007/3-540-16442-1_14.

[7] M. Herlihy & N. Shavit (2008): *The Art of Multiprocessor Programming*. Morgan Kaufmann.

[8] M. P. Herlihy & J. M. Wing (1990): *Linearizability: a correctness condition for concurrent objects*. *ACM Trans. Program. Lang. Syst.* 12(3), pp. 463–492, doi:10.1145/78969.78972.

[9] L. Lamport (1979): *How to Make a Correct Multiprocess Program Execute Correctly on a Multiprocessor*. *IEEE Trans. Computers* 46(7), pp. 779–782, doi:10.1109/12.599898.

[10] R. K. Treiber (1986): *Systems programming: Coping with parallelism*. Technical Report RJ 5118, IBM Almaden Res. Ctr.